TM-537
TM-537
2340.000
December 1974

BISON-NET, A HIGH SPEED BLOCK MULTIPLEXER CHANNEL FOR INTERCONNECTING COMPUTERS VIA CAMAC

A. E. Brenner, Thomas F. Droege and Ronald G. Martin

Fermi National Accelerator Laboratory*
P. O. Box 500
Batavia, Illinois 60510

Summary

A high speed network communications station has been developed to serve to interconnect a variety of computers in a multiplexed star network. Each station consists of one control module and two memory modules embodied as CAMAC modules. Transmission is over coaxial lines with typical distances between stations of from 1 to 2 miles. The transmission rate is 4MHz.

Introduction

At Fermi National Accelerator Laboratory typically a number of experiments are simultaneously set up and working in each of five geographically separated experimental areas. Figure 1 illustrates the approximate Laboratory layout. The accelerator cycle typically runs about three to six seconds of acceleration period during which time little of interest to the experimenters occurs and then a constant energy period during which time the beam is extracted slowly over a period of that one second to most experiments. During this one second spill time, based upon high speed electronics logic decisions, an interrupt to a local on-line computer system is generated and data is strobed from the high speed electronics into the computer. This data strobe typically is of the order of 100 mini-computer words of data and frequently of the order of 100 triggers occur during one spill period. The mini-computers collect data in this manner, buffer and frequently format the data and then store this data on tape as the primary data long term storage medium. The mini's also operate on some fraction of the data to monitor the operation of the experiment, usually in close interaction with the experimenter.

Frequently it occurs, however, that the mini-computer is inadequate in computing power to give an adequate physics result to the experimenter. For this reason it is necessary to have additional computing power available to make these more extensive calculations. Given today's technology it is possible to add floating point hardware to a modern mini-computer to alleviate this situation. However, since there are a large number of mini-computers installed in our Laboratory, and also because frequently more computing power is required than can be obtained with typical mini-computer floating point hardware, it is more appropriate to connect them with a real time link to a yet larger more powerful computer. At Fermilab we use a CDC 6600 in this role. At the present time however, data is brought to this computer by bringing the magnetic tape to the computer via the "on-line bicycle route". It is to remove this rather slow "on-line" mechanism that we have designed a general purpose Laboratory wide data communications network to link our mini-computers to a central large computer. This network is called BISON-NET.

Global Concepts

Our design for BISON-NET is based upon the following considerations:

a) There are a variety of mini-computers of different manufacture servicing the various experiments.

b) Almost all the mini-computers use CAMAC as the normal interface system.

c) Although the CAMAC protocol is rather cumbersome for the major data acquisition and frequently tends to be the limiting element during data strobes, nevertheless during the acceleration cycle the activity level for the CAMAC branch highway is very low.

d) The network should handle reasonable sized buffers from each experimenter, with an average peak rate of one buffer per experimenter per acceleration cycle guaranteed.

e) As much as possible the communications protocol should be transparent to the user.

f) There should be no distance limit between terminals.

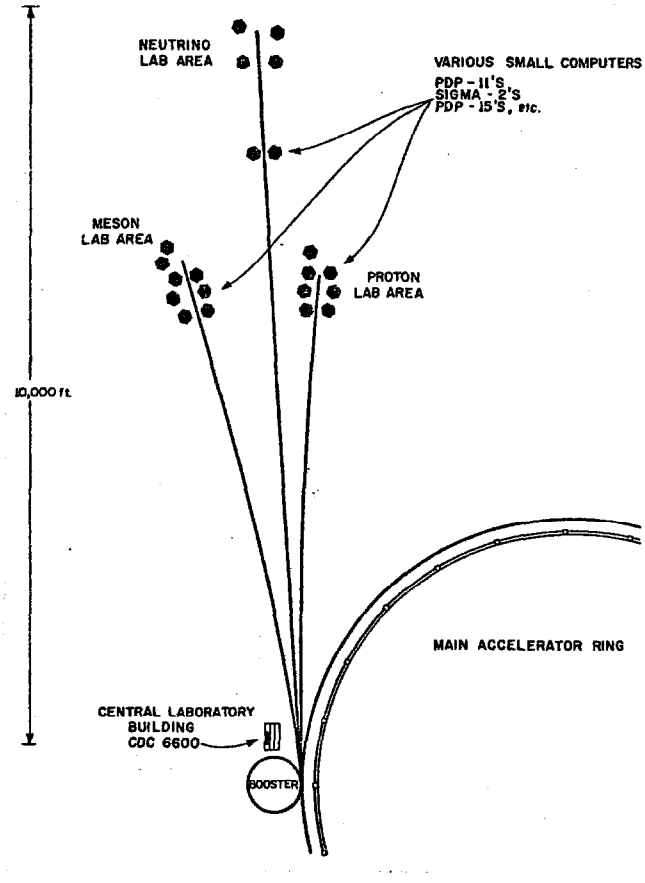

Fig. 1 Fermilab Experimental Area Layout

These considerations lead to a design which utilizes a CAMAC interfaced system. The implementation which is symmetric at transmit and receive ends requires only two different types of modules. As illustrated in Figure 2, a network terminal requires one transmit/receive control module (T/R Module) and two 1024 word 24 bit memory modules, one serving as a transmit memory (TM Module) and one as a receiving memory (RM Module). A typical mini-computer to CDC

*Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, Illinois 60510, Operated by Universities Research Association Inc. Under Contract with U.S. AEC

6600 connection is given in Figure 3. The whole network is a star network and because of the rather large distances involved and the relatively high cost of high quality transmission cable, multiplexers are utilized for physically close clusters. The network is schematically illustrated in Figure 4.

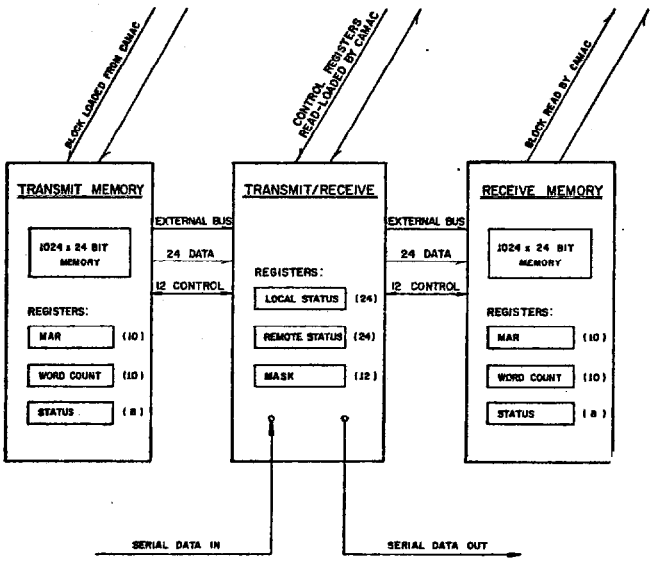

Fig. 2 Major Features of BISON-NET CAMAC Modules.

Typically at the end of a spill period, when the CAMAC system is not in a time critical situation, the user initiates a block transfer from a 1K buffer in his core memory into the 1K buffer BISON-NET TM module. Up to a dozen different users might do this more or less simultaneously. Once the transmit memory is loaded the transmit/receive module transmits the whole block serially to the receive memory at the CDC end. The multiplexers take care of the priority and protocol if two TM modules become filled at the same moment. At the receive end appropriate status bits are set and the software at the 6600 end scans the status registers and unloads the data in an appropriate fashion. The T/R modules are reasonably independent. They are designed to take care of a large variety of unusual error conditions and will automatically retransmit the block if there has been an unsuccessful transmission.

The data transmitted via this technique is either experimental data or control words. The first four words are software set and define the details of the accompanying data. Typically the user is allowed to construct a file which is controlled by special software in the CDC 6600 peripheral and central processors. These are buffered in core memory in the CDC 6600 and written into the disk by the system software. The user may open and close this file and after the file is closed, he may submit high priority batch jobs by transmitting appropriate control words and associated control blocks. Normally it is assumed that the data files are constructed over a period of minutes or tens of minutes, and at the end of this time, a high priority batch job is initiated. This batch program may construct an output file if appropriate and upon initiation by the mini-computer, this block will be transmitted back to the mini-computer.

Memory Module

One of two types of CAMAC modules required for the BISON-NET is the 1024 word by 24 bit memory (Figure 2). This module uses static N-MOS memory, TTL logic, and has been designed as a general purpose memory. Connection to the T/R module is via a rear panel connector located above the CAMAC dataway connector. Twenty four lines are used as a bi-directional dataway (the T/R dataway) with the remainder used for control functions. Memory modules function either as a transmitting (TM) or a receiving memory (RM) depending on where their interconnecting cable is attached to the T/R module.

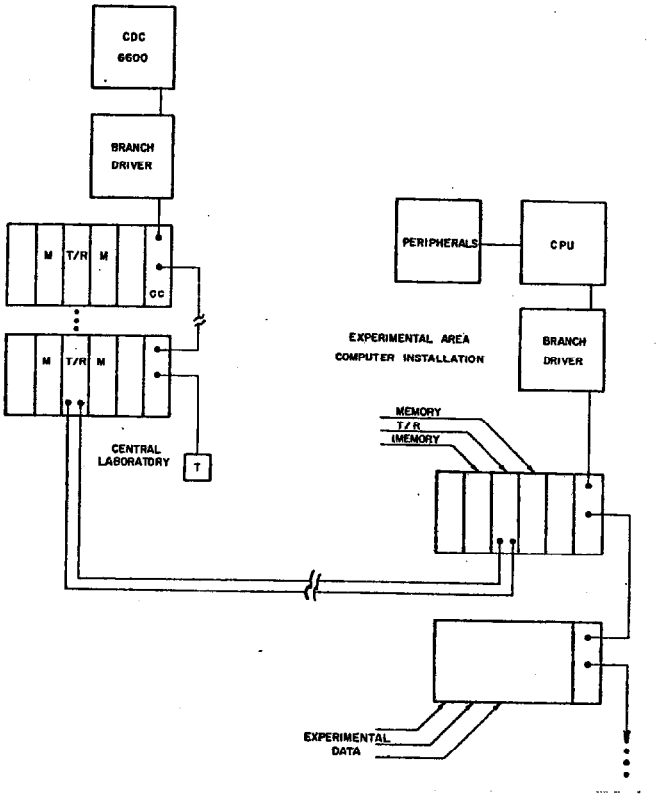

Fig. 3 Typical Mini-computer to CDC 6600 Connection

To aid in system checkout and error diagnosis, all important control registers can be set and cleared by CAMAC commands. In addition various sections of the automatic logic can be inhibited by setting control bits. The CAMAC commands and register configuration are detailed in Table 1.

A normal BISON-NET block transfer operation starts with an F(16) block write command to the TM module. The first word of a block transfer is treated in special way by the TM module. The low order ten bits of the first word give the word count of the block transfer, and are set by software. The value of this word count is stored in word count register (WCR), and is continuously compared with the memory address register (MAR) during transfer. The first word is also stored in its entirety in the first location of the MOS memory. The MAR is incremented at the CAMAC S2 time and successive words are stored in successive memory locations. The TM module monitors the rate at which words are received from CAMAC and when a break is determined by a pause greater than a preset time (as determined by the particular computer branch driver combination) the memory load operation is terminated. If the MAR and the WCR do not match, an error condition is indicated. The MAR is then reset to zero.

If this match is successful and if no other

error conditions exist, the T/R module receives a signal indicating that a block is available for transfer to the remote station. The T/R module now switches the memory data path from reading from the CAMAC dataway to writing on the T/R dataway. The T/R module serializes the memory words, adds parity and control bits, generates Di-phase code (similar to Manchester code) for transmission and transmits the serial data at its internal clock rate. The MAR is advanced by pseudo S2 pulses generated in the T/R module and transmission continues until terminated by equality between the MAR and WCR.

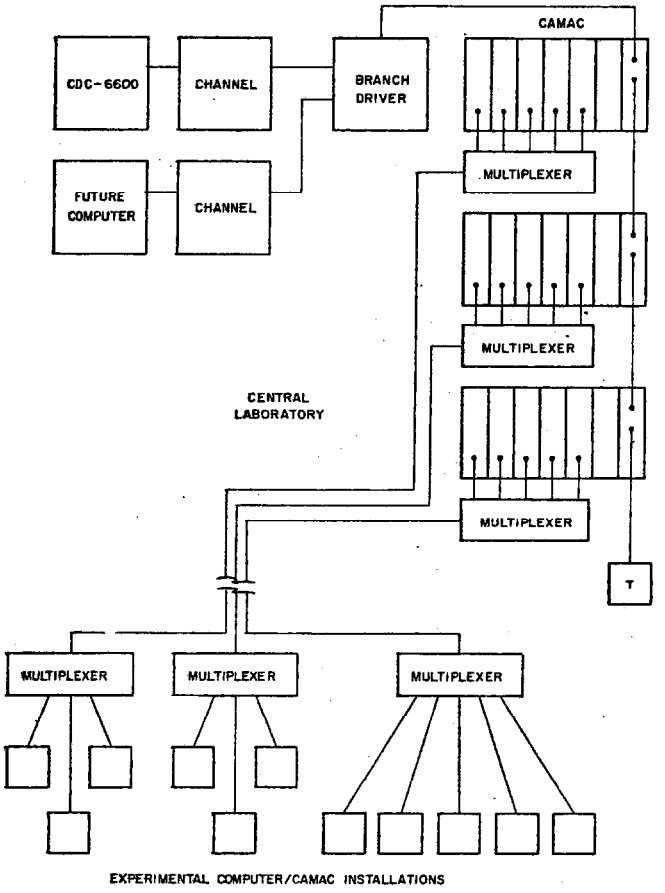

Fig. 4 BISON-NET network with multiplexers shown which allow shared use of cables between the Central Laboratory and the principal experimental areas.

Operational failures are detected by the T/R module which sets a LAM to indicate that something has gone wrong in a memory module. A memory status register can then be examined to determine the exact cause of the error. Note that T/R module LAMs are reserved for error conditions and memory LAMs for successful completions.

A number of error conditions are possible. Blocks of improper length, attempts to write into a RM module or to read from a TM module or to write blocks longer than 1024 words are all detected. A T/R module LAM will also be set if a write is attempted to a TM module which has not as yet successfully transferred its contents to the remote location or if a read is attempted from a RM module that has not received new data. It is thus possible to program blindly without examining status registers or servicing LAMs if it is known ahead of time that the probability of a successful operation is high. Thus an experiment passing one buffer per beam spill could write blindly with the expectation that the last buffer would almost always be clear. In the rare instances where this is not the case (e.g., the central computer is down briefly) it is only necessary to attempt to reload the TM module since logic prevents overwriting data which has not been successfully transmitted.

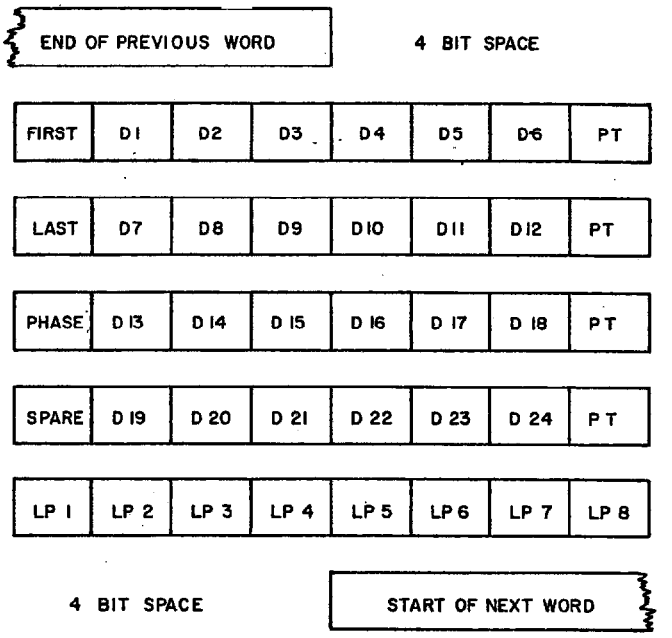

Fig. 5 BISON-NET word format, PT is odd transverse parity while LPn is even column parity.

Similar logic holds for the RM module. In this case the T/R module fabricates pseudo S2 pulses, assembles full words and then transmits the input data to the RM module via the T/R dataway. When transmission errors are detected by the parity circuits, RM module load operations are terminated. This results in a time out with an improper length comparison and the memory automatically resets its address register to zero to look for a retransmission. Except for special modes of operation, the memory address register is reset to zero at each timeout so that the memory is always ready to load a block starting at location zero.

Memory LAMs are set only on successful completion of a memory unload operation. Thus the TM module sets a LAM when its contents are successfully transferred to the remote location and the RM module sets a LAM when its contents have been unloaded via CAMAC.

T/R Module

The T/R modules at each end of the transmission link control the traffic between the four memory modules. Operation is almost full duplex in that transmission can occur simultaneously in both directions. A key feature of the T/R module (Figure 2) is its two status registers. Each T/R module contains a 24 bit local status register (LSR) which indicates all its important control functions. In addition each T/R module also contains a second register, the remote status register (RSR), which holds the most recently received contents of the local status register from the T/R module at the remote station. Thus each T/R module is constantly aware of its own status and the status of the remote T/R module. If there is no other traffic, status is transmitted at about 0.1 second

intervals. One of the LSR bits, and a front panel light indicate that the T/R module is regularly receiving status from the remote station. Since this equivalent bit from the remote T/R module is returned in the RSR and is indicated with a second light, quick visual determination that transmission is taking place in both directions is possible. If this status protocol fails a T/R module LAM is set. The CAMAC commands and register identification for the T/R module are listed in Table 2.

In general, any change of the LSR initiates the transmission of its contents. Since this is often critical information, status words have priority over data words and are inserted in a data block whenever they occur. Two bits in each transmitted word are used to identify the first word in a data block, a word in a block, the last word in a data block, and a status word.

TABLE 1

MEMORY MODULE

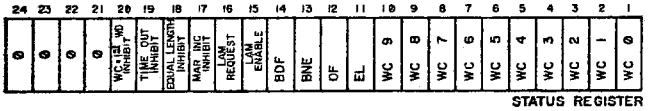

STATUS REGISTER

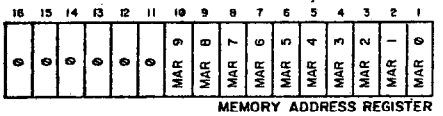

MEMORY ADDRESS REGISTER

CAMAC COMMANDS	FUNCTION
F(0) · A(n)	MEMORY BLOCK READ
F(1) · A(0)	READ WC & STATUS REGISTER
F(1) · A(1)	READ SINGLE MEMORY LOCATION [AS f(MAR)]
F(1) · A(2)	READ MAR
F(2) · A(n)	OPTIONAL BLOCK READ [INSTEAD OF F(0)]
F(8) · A(0) · S1	TEST LAM Q=1 IF LAM REQ. · MASK
F(10) · A(0) · S1	RESET LAM REQ.
F(11) · A(0) · S1	CLEAR WC & STATUS REGISTER
F(11) · A(2) · S1	CLEAR MAR
F(16) · A(n) · S1	MEMORY BLOCK WRITE
F(17) · A(1) · S1	LOAD SINGLE MEMORY LOCATION [AS f(MAR)]
F(17) · A(2) · S1	LOAD MAR
F(18) · A(n) · S1	OPTIONAL BLOCK WRITE [INSTEAD OF F(16)]
F(19) · A(0) · S1	SELECTIVE SET WC & STATUS REGISTER
F(24) · A(0) · S1	DISABLE LAM
F(26) · A(0) · S1	ENABLE LAM
F(27) · A(0) · S1	TEST LAM Q=1 IF LAM REQ.
(Z + C) · S2	CLEAR ALL REGISTERS

NOTE: "X" RESPONSE GENERATED FOR ALL VALID COMMANDS AND "Q" RESPONSE GENERATED FOR READ AND WRITE OPERATIONS

Transmission is by 40 bit data words separated by 4 bit spaces. The data format (Figure 5) consists of 24 data bits, 4 lateral (odd) parity bits, two word identifying bits, a phase bit, and 8 long (even) parity bits. An extra bit allows later system expansion. Since the receiver circuitry is time sensitive, requiring 40 data bits at regular intervals followed by a space, the transmission link is less sensitive to white noise than the 12 parity bits would otherwise indicate. It is expected that this scheme is sufficiently protective that the undetected error rate will be insignificant.[2] In the design, emphasis has been placed on error detection and retransmission rather than error correction.

As mentioned earlier, T/R module status bits prevent loading a TM module until its previous contents have been successfully transferred to the remote station. Once the TM module is loaded, the T/R attempts to transfer the memory contents. First the T/R checks to see if the last block transmitted has been unloaded. If not, the T/R waits indefinitely for a status condition indicating the unload has taken place. Initially this wait was to be terminated by a short hardware time out, however global considerations make it desirable to allow a rather long wait to accommodate unusual load conditions at the central computer.

TABLE 2

T/R MODULE

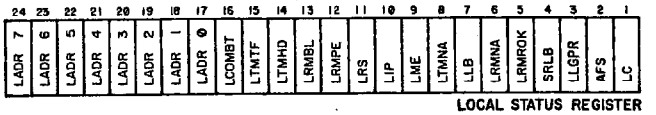

LOCAL STATUS REGISTER

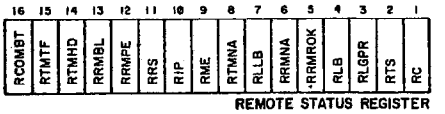

REMOTE STATUS REGISTER

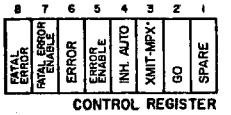

CONTROL REGISTER

CAMAC COMMANDS	FUNCTION
F(1) · A(0)	READ LOCAL STATUS REGISTER
F(1) · A(1)	READ REMOTE STATUS REGISTER
F(1) · A(2)	READ CONTROL REGISTER
F(8) · A(0) · S1	TEST FATAL ERROR LAM Q=1 IF SET
F(10) · A(0) · S1	RESET FATAL ERROR
F(11) · A(0) · S1	CLEAR LOCAL STATUS REGISTER
F(11) · A(1) S1	CLEAR REMOTE STATUS REGISTER
F(11) · A(2) · S1	CLEAR CONTROL REGISTER
F(11) · A(3) · S1	CLEAR ALL REGISTERS
F(17) · A(1) · S1	WRITE LOCAL STATUS COMMUNICATION BIT
F(19) · A(0) · S1	SELECTIVE SET LOCAL STATUS REGISTER
F(19) · A(1) · S1	WRITE REMOTE STATUS REGISTER
F(19) · A(2) S1	SELECTIVE SET CONTROL REGISTER
(Z + C) · S2	CLEAR ALL REGISTERS

NOTE: "X" RESPONSE GENERATED FOR ALL VALID COMMANDS AND "Q" RESPONSE GENERATED FOR READ AND WRITE OPERATIONS

Once the T/R module determines that the remote memory is empty, transmission is attempted. The T/R module then waits somewhat longer than the round trip cable time for a status word acknowledging that the block has been received correctly. If this is not received the block is retransmitted. This continues until transmission is successful or a hardware time out occurs.

Since only the data blocks are acknowledged, loss of a status word could result in the retransmission of a correctly received block. If the receiving station computer happened to be very prompt in unloading its memory it could receive the same block twice. To prevent this, each word contains a phase bit which alternates on successive blocks. Comparison of the phase of an incoming block with the phase of the last good block received allows discarding of repeat transmissions of good blocks.

In order to allow shared use of transmission facilities, the first word of each block and all status words contain an 8 bit T/R module address which is automatically inserted by the hardware in the high order 8 bits. In order for communication to take place the two T/R modules must be selecter switch set to the same address.

The T/R module contains a number of features which aid in the detection and correction of troubles. First and foremost is the continuous exchange of status which constantly exercises about 90% of the hardware. If this fails, the local T/R module can be tested by setting the loop back bit of the local status register which logically connects the output to the input line and tests everything but the line drivers. Setting the remote loop back bit in the remote module via status transmission requires only a fraction of the remote module to be operational and tests the transmission line both ways. Other control bits allow set-

ting a LAM to interrupt the remote module, and immediate status transmission from the local and remote module. Automatic operation can be inhibited and transmission can be controlled step by step using the status registers and various interrupts.

The Di-phase code used for transmission is self clocking, requiring only approximate match of the receiving module clock to the transmitted data and only one transmission line in each direction.[1] The transmission line (presently coaxial cable) is driven by a balanced transformer coupled driver for good common mode rejection.

Multiplexer

Since each transmitted message carries the address of the pair of the T/R modules, this is used to allow several sets of modules to use the same transmission link. The scheme used is to poll the multiplexer T/R modules until one is found with a message. This module is then given exclusive use of the link until it has completed transmission of a single block or status word. If retransmission is required, however, it is placed at the end of a queue. At the receiving end, all modules receive all messages in parallel, ignoring messages addressed to other modules.

Current Status

Three T/R modules and five memory modules have been fabricated and tested. This is enough for one complete link with a spare for each position. These modules have been operated successfully at a 4 MHZ bit rate between two PDP-11's connected by 11,000 feet of cable, system checkout is nearly complete with presently planned module logic changes being required only to aid system programming or diagnostic software. The T/R module and the memory module require approximately 180 and 160 integrated circuits respectively.

The operating system software is completely specified with coding progressing for both the PDP-11 user software and the CDC 6600 system software. Individual module diagnostic routines have been written as well as a very flexible CAMAC Macro diagnostic language[3] which is usable for any PDP-11 CAMAC system using a BD0-11.

A contract has been placed for a Branch Driver for the CDC 6600. The vendor is presently testing their hardware connection to the CDC 6600 channel.

Operation on line with an experiment is expected by the end of the first quarter of 1975 with a dozen or so users by mid-year.

Acknowledgements

We wish to thank L. Leipuner and K. Rich for contributions to the system design philosophy and also D. Curtis and J. Mack for producing the diagnostic software. M. Watson and T. Wesson contributed many design details and skillfully assembled the prototypes.

References

1. Signetis Application Memo 8T20
2. R. K. Richards, Digital Design, John Wiley & Sons, New York, New York, Chapter 5.
3. J. Mack, "CAMAC Diagnostic Program", Fermilab Computer Department Internal Report PN-55, Fermi National Accelerator Laboratory, Batavia, Illinois